arXiv:1703.02901v1 [cs.CG] 8 Mar 2017

Local Equivalence and Intrinsic Metrics between Reeb Graphs *

Mathieu Carrière, Steve Oudot

March 9, 2017

Abstract

As graphical summaries for topological spaces and maps, Reeb graphs are common objects in the computer graphics or topological data analysis literature. Defining good metrics between these objects has become an important question for applications, where it matters to quantify the extent by which two given Reeb graphs differ. Recent contributions emphasize this aspect, proposing novel distances such as *functional distortion* or *interleaving* that are provably more discriminative than the so-called *bottleneck distance*, being true metrics whereas the latter is only a pseudo-metric. Their main drawback compared to the bottleneck distance is to be comparatively hard (if at all possible) to evaluate. Here we take the opposite view on the problem and show that the bottleneck distance is in fact good enough *locally*, in the sense that it is able to discriminate a Reeb graph from any other Reeb graph in a small enough neighborhood, as efficiently as the other metrics do. This suggests considering the *intrinsic metrics* induced by these distances, which turn out to be all *globally* equivalent. This novel viewpoint on the study of Reeb graphs has a potential impact on applications, where one may not only be interested in discriminating between data but also in interpolating between them.

1 Introduction

In the context of shape analysis, the Reeb graph [25] provides a meaningful summary of a topological space and a real-valued function defined on that space. Intuitively, it continuously collapses the connected components of the level sets of the function into single points, thus tracking the values of the functions at which the connected components merge or split. Reeb graphs have been widely used in computer graphics and visualization—see [6] for a survey, and their discrete versions, including the so-called *Mappers* [26], have become emblematic tools of topological data analysis due to their success in applications [2, 3, 19, 22].

Finding relevant dissimilarity measures for comparing Reeb graphs has become an important question in the recent years. The quality of a dissimilarity measure is usually assessed through three criteria: its ability to satisfy the axioms of a metric, its discriminative power, and its computational efficiency. The most natural choice to begin with is to use the *Gromov-Hausdorff distance* d_{GH} [9] for Reeb graphs seen as metric spaces. The main drawback of this distance is to quickly become intractable to compute in practice, even for graphs that are metric trees [1]. Among recent contributions, the *functional distortion distance* d_{FD} [4] and the *interleaving distance* d_{I} [14] share

*This work was partially supported by ERC Grant Agreement No. 339025 GUDHI (Algorithmic Foundations of Geometry Understanding in Higher Dimensions) and was carried out while the second author was visiting the ICERM at Brown University.

the same advantages and drawbacks as d_{GH}, in particular they enjoy good stability and discriminativity properties but they lack efficient algorithms for their computation, moreover they can be difficult to interpret. By contrast, the *bottleneck distance* d_{B} comes with a signature for Reeb graphs, called the *extended persistence diagram* [13], which acts as a stable bag-of-feature descriptor. Furthermore, d_{B} can be computed efficiently in practice. Its main drawback though is to be only a pseudo-metric, so distinct graphs can have the same signature and therefore be deemed equal in d_{B}.

Another desired property for dissimilarity measures is to be *intrinsic*, i.e. realized as the lengths of shortest continuous paths in the space of Reeb graphs [9]. This is particularly useful when one actually needs to interpolate between data, and not just discriminate between them, which happens in applications such as image or 3-d shape morphing, skeletonization, and matching [17, 20, 21, 27]. At this time, it is unclear whether the metrics proposed so far for Reeb graphs are intrinsic or not. Using intrinsic metrics would not only open the door to the use of Reeb graphs in the aforementioned applications, but it would also provide a better understanding of the intrinsic structure of the space of Reeb graphs, and give a deeper meaning to the distance values.

Our contributions. In the first part of the paper we show that the bottleneck distance can discriminate a Reeb graph from any other Reeb graph in a small enough neighborhood, as efficiently as the other metrics do, even though it is only a pseudo-metric globally. More precisely, we show that, given any constant $K \in (0, 1/22]$, in a sufficiently small neighborhood of a given Reeb graph R_f in the functional distortion distance (that is: for any Reeb graph R_g such that $d_{\mathrm{FD}}(\mathrm{R}_f, \mathrm{R}_g) < c(f, K)$, where $c(f, K) > 0$ is a positive constant depending only on f and K), one has:

$$K d_{\mathrm{FD}}(\mathrm{R}_f, \mathrm{R}_g) \leq d_{\mathrm{B}}(\mathrm{R}_f, \mathrm{R}_g) \leq 2 d_{\mathrm{FD}}(\mathrm{R}_f, \mathrm{R}_g). \tag{1}$$

The second inequality is already known [4], and it asserts that the bottleneck distance between Reeb graphs is stable. The first inequality is new, and it asserts that the bottleneck distance is discriminative locally, in fact just as discriminative as the other distances mentioned above. Equation (1) can be viewed as a local equivalence between metrics although not in the usual sense: firstly, all comparisons are anchored to a fixed Reeb graph R_f, and secondly, the constants K and 2 are absolute.

The second part of the paper advocates the study of intrinsic metrics on the space of Reeb graphs, for the reasons mentioned above. As a first step, we propose to study the intrinsic metrics $\hat{d}_{\mathrm{GH}}$, $\hat{d}_{\mathrm{FD}}$, $\hat{d}_{\mathrm{I}}$ and $\hat{d}_{\mathrm{B}}$ induced respectively by d_{GH}, d_{FD}, d_{I} and d_{B}. While the first three are obviously globally equivalent because their originating metrics are, our second contribution is to show that the last one is also globally equivalent to the other three.

The paper concludes with a discussion and some directions for the study of the space of Reeb graphs as an intrinsic metric space.

Related work. Interpolation between Reeb graphs is also the underlying idea of the *edit distance* recently proposed by Di Fabio and Landi [15]. The problem with this distance, in its current form at least, is that it restricts the interpolation to pairs of graphs lying in the same homeomorphism class. By contrast, our class of admissible paths is defined with respect to the topology induced by the functional distortion distance, as such it allows interpolating between distinct homeomorphism classes.

Interpolation between Reeb graphs is also related to the study of inverse problems in topological data analysis. To our knowledge, the only result in this vein shows the differentiability of the operator sending point clouds to the persistence diagram of their distance function [16]. Our first contribution (1) sheds light on the operator's local injectivity properties over the class of Reeb graphs.

2 Background

Throughout the paper we work with singular homology with coefficients in the field $\mathbb{Z}_2$, which we omit in our notations for simplicity. In the following, "connected" stands for "path-connected", and "cc" stands for "connected component(s)". Given a map $f : X \to \mathbb{R}$ and an interval $I \subseteq \mathbb{R}$, we write X_f^I as a shorthand for the preimage $f^{-1}(I)$, and we omit the subscript when the map is obvious from the context.

2.1 Morse-Type Functions

Definition 2.1. *A continuous real-valued function f on a topological space X is* of Morse type *if:*

(i) there is a finite set $Crit(f) = \{a_1 < ... < a_n\} \subset \mathbb{R}$, called the set of critical values*, such that over every open interval $(a_0 = -\infty, a_1), ..., (a_i, a_{i+1}), ..., (a_n, a_{n+1} = +\infty)$ there is a compact and locally connected space Y_i and a homeomorphism $\mu_i : Y_i \times (a_i, a_{i+1}) \to X^{(a_i,a_{i+1})}$ such that $\forall i = 0, ..., n, f|_{X^{(a_i,a_{i+1})}} = \pi_2 \circ \mu_i^{-1}$, where π_2 is the projection onto the second factor;*

(ii) $\forall i = 1, ..., n-1, \mu_i$ extends to a continuous function $\bar{\mu}_i : Y_i \times [a_i, a_{i+1}] \to X^{[a_i,a_{i+1}]}$; similarly, μ_0 extends to $\bar{\mu}_0 : Y_0 \times (-\infty, a_1] \to X^{(-\infty,a_1]}$ and μ_n extends to $\bar{\mu}_n : Y_n \times [a_n, +\infty) \to X^{[a_n,+\infty)}$;

(iii) Each levelset $f^{-1}(t)$ has a finitely-generated homology.

Let us point out that a Morse function is also of Morse type, and that its critical values remain critical in the definition above. Note that some of its regular values may be termed critical as well in this terminology, with no effect on the analysis.

2.2 Extended Persistence

Let f be a real-valued function on a topological space X. The family $\{X^{(-\infty,\alpha]}\}_{\alpha\in\mathbb{R}}$ of sublevel sets of f defines a *filtration*, that is, it is nested w.r.t. inclusion: $X^{(-\infty,\alpha]} \subseteq X^{(-\infty,\beta]}$ for all $\alpha \leq \beta \in \mathbb{R}$. The family $\{X^{[\alpha,+\infty)}\}_{\alpha\in\mathbb{R}}$ of superlevel sets of f is also nested but in the opposite direction: $X^{[\alpha,+\infty)} \supseteq X^{[\beta,+\infty)}$ for all $\alpha \leq \beta \in \mathbb{R}$. We can turn it into a filtration by reversing the order on the real line. Specifically, let $\mathbb{R}^{\mathrm{op}} = \{\tilde{x} \mid x \in \mathbb{R}\}$, ordered by $\tilde{x} \leq \tilde{y} \Leftrightarrow x \geq y$. We index the family of superlevel sets by $\mathbb{R}^{\mathrm{op}}$, so now we have a filtration: $\{X^{[\tilde{\alpha},+\infty)}\}_{\tilde{\alpha}\in\mathbb{R}^{\mathrm{op}}}$, with $X^{[\tilde{\alpha},+\infty)} \subseteq X^{[\tilde{\beta},+\infty)}$ for all $\tilde{\alpha} \leq \tilde{\beta} \in \mathbb{R}^{\mathrm{op}}$.

Extended persistence connects the two filtrations at infinity as follows. First, replace each superlevel set $X^{[\tilde{\alpha},+\infty)}$ by the pair of spaces $(X, X^{[\tilde{\alpha},+\infty)})$ in the second filtration. This maintains the filtration property since we have $(X, X^{[\tilde{\alpha},+\infty)}) \subseteq (X, X^{[\tilde{\beta},+\infty)})$ for all $\tilde{\alpha} \leq \tilde{\beta} \in \mathbb{R}^{\mathrm{op}}$. Then, let $\mathbb{R}_{\mathrm{Ext}} = \mathbb{R} \cup \{+\infty\} \cup \mathbb{R}^{\mathrm{op}}$, where the order is completed by $\alpha < +\infty < \tilde{\beta}$ for all $\alpha \in \mathbb{R}$ and $\tilde{\beta} \in \mathbb{R}^{\mathrm{op}}$. This poset is isomorphic to $(\mathbb{R}, \leq)$. Finally, define the *extended filtration* of f over $\mathbb{R}_{\mathrm{Ext}}$ by:

$$F_\alpha = X^{(-\infty,\alpha]} \text{ for } \alpha \in \mathbb{R},\ F_{+\infty} = X \equiv (X, \emptyset) \text{ and } F_{\tilde{\alpha}} = (X, X^{[\tilde{\alpha},+\infty)}) \text{ for } \tilde{\alpha} \in \mathbb{R}^{\mathrm{op}},$$

where we have identified the space X with the pair of spaces $(X, \emptyset)$ at infinity. The subfamily $\{F_\alpha\}_{\alpha\in\mathbb{R}}$ is the *ordinary* part of the filtration, while $\{F_{\tilde{\alpha}}\}_{\tilde{\alpha}\in\mathbb{R}^{\mathrm{op}}}$ is the *relative* part.

Applying the homology functor H_* to this filtration gives the so-called *extended persistence module* $\mathbb{V}$ of f, which is a sequence of vector spaces connected by linear maps induced by the inclusions in the extended filtration. For functions of Morse type, the extended persistence module can be decomposed as a finite direct sum of half-open *interval modules*—see e.g. [12]: $\mathbb{V} \simeq \bigoplus_{k=1}^{n} \mathbb{I}[b_k, d_k)$, where each summand $\mathbb{I}[b_k, d_k)$ is made of copies of the field of coefficients at every index $\alpha \in [b_k, d_k)$, and of copies of the zero space elsewhere, the maps between copies of the field being identities. Each summand represents the lifespan of a *homological feature* (cc, hole, void, etc.) within the filtration. More precisely, the *birth time* b_k and *death time* d_k of the feature are given by the endpoints of the interval. Then, a convenient way to represent the structure of the module is to plot each interval in the decomposition as a point in the extended plane, whose coordinates are given by the endpoints. Such a plot is called the *extended persistence diagram* of f, denoted $\mathrm{Dg}(f)$. The distinction between ordinary and relative parts of the filtration allows us to classify the points in $\mathrm{Dg}(f)$ as follows:

- $p = (x, y)$ is called an *ordinary* point if $x, y \in \mathbb{R}$;
- $p = (x, y)$ is called a *relative* point if $x, y \in \mathbb{R}^{\mathrm{op}}$;
- $p = (x, y)$ is called an *extended* point if $x \in \mathbb{R}, y \in \mathbb{R}^{\mathrm{op}}$;

Note that ordinary points lie strictly above the diagonal $\Delta = \{(x, x) \mid x \in \mathbb{R}\}$ and relative points lie strictly below Δ, while extended points can be located anywhere, including on Δ (e.g. when a cc lies inside a single critical level, see Section 2.3). It is common to partition $\mathrm{Dg}(f)$ according to this classification: $\mathrm{Dg}(f) = \mathrm{Ord}(f) \sqcup \mathrm{Rel}(f) \sqcup \mathrm{Ext}^+(f) \sqcup \mathrm{Ext}^-(f)$, where by convention $\mathrm{Ext}^+(f)$ includes the extended points located on the diagonal Δ.

Stability. An important property of extended persistence diagrams is to be stable in the so-called *bottleneck distance* d_b^∞. Given two persistence diagrams D, D', a *partial matching* between D and D' is a subset Γ of $D \times D'$ where for every $p \in D$ there is at most one $p' \in D'$ such that $(p, p') \in \Gamma$, and conversely, for every $p' \in D'$ there is at most one $p \in D$ such that $(p, p') \in \Gamma$. Furthermore, Γ must match points of the same type (ordinary, relative, extended) and of the same homological dimension only. The *cost* of Γ is: $\mathrm{cost}(\Gamma) = \max\{\max_{p\in D} \delta_D(p), \max_{p'\in D'} \delta_{D'}(p')\}$, where $\delta_D(p) = \|p - p'\|_\infty$ if p is matched to some $p' \in D'$ and $\delta_D(p) = d_\infty(p, \Delta)$ if p is unmatched—same for $\delta_{D'}(p')$.

Definition 2.2. *The* bottleneck distance *between two persistence diagrams D and D' is $d_{\mathrm{B}}(D, D') = \inf_\Gamma cost(\Gamma)$, where Γ ranges over all partial matchings between D and D'.*

Theorem 2.3 (Stability [13])**.** *For any Morse-type functions $f, g : X \to \mathbb{R}$,*

$$d_{\mathrm{B}}(\mathrm{Dg}(f), \mathrm{Dg}(g)) \leq \|f - g\|_\infty. \tag{2}$$

2.3 Reeb Graphs

Definition 2.4. *Given a topological space X and a continuous function $f : X \to \mathbb{R}$, we define the equivalence relation $\sim_f$ between points of X by $x \sim_f y$ if and only if $f(x) = f(y)$ and x, y belong to the same cc of $f^{-1}(f(x)) = f^{-1}(f(y))$. The* Reeb graph $\mathrm{R}_f(X)$ *is the quotient space $X/\sim_f$. As f is constant on equivalence classes, there is a well-defined induced map $\tilde{f} : \mathrm{R}_f(X) \to \mathbb{R}$.*

Connection to extended persistence. If f is a function of Morse type, then the pair (X, f) is an $\mathbb{R}$-*constructible space* in the sense of [14]. This ensures that the Reeb graph is a multigraph, whose nodes are in one-to-one correspondence with the cc of the critical level sets of f. In that case, there is a nice interpretation of $\mathrm{Dg}(\tilde{f})$ in terms of the structure of $\mathrm{R}_f(X)$. We refer the reader to [4, 13] and the references therein for a full description as well as formal definitions and statements. Orienting the Reeb graph vertically so $\tilde{f}$ is the height function, we can see each cc of the graph as a trunk with multiple branches (some oriented upwards, others oriented downwards) and holes. Then, one has the following correspondences, where the *vertical span* of a feature is the span of its image by $\tilde{f}$:

- The vertical spans of the trunks are given by the points in $\mathrm{Ext}_0^+(\tilde{f})$;
- The vertical spans of the downward branches are given by the points in $\mathrm{Ord}_0(\tilde{f})$;
- The vertical spans of the upward branches are given by the points in $\mathrm{Rel}_1(\tilde{f})$;
- The vertical spans of the holes are given by the points in $\mathrm{Ext}_1^-(\tilde{f})$.

The rest of the diagram of $\tilde{f}$ is empty. These correspondences provide a dictionary to read off the structure of the Reeb graph from the persistence diagram of the quotient map $\tilde{f}$. Note that it is a bag-of-features type of descriptor, taking an inventory of all the features together with their vertical spans, but leaving aside the actual layout of the features. As a consequence, it is an incomplete descriptor: two Reeb graphs with the same persistence diagram may not be isomorphic. See the two Reeb graphs in Figure 1 for instance.

Notation. Throughout the paper, we consider Reeb graphs coming from Morse-type functions, equipped with their induced maps. We denote by Reeb the space of such graphs. In the following, we have $\mathrm{R}_f, \mathrm{R}_g \in \mathsf{Reeb}$, with induced maps $f : \mathrm{R}_f \to \mathbb{R}$ with critical values $\{a_1, ..., a_n\}$, and $g : \mathrm{R}_g \to \mathbb{R}$ with critical values $\{b_1, ..., b_m\}$. Note that we write f, g instead of $\tilde{f}, \tilde{g}$ for convenience. We also assume without loss of generality (w.l.o.g.) that R_f and R_g are connected. If they are not connected, then our analysis can be applied component-wise.

2.4 Distances for Reeb graphs

Definition 2.5. *The* bottleneck distance *between* R_f *and* R_g *is:*

$$d_{\mathrm{B}}(\mathrm{R}_f, \mathrm{R}_g) := d_{\mathrm{B}}(\mathrm{Dg}(f), \mathrm{Dg}(g)). \tag{3}$$

Definition 2.6. *The* functional distortion distance *between* R_f *and* R_g *is:*

$$d_{\mathrm{FD}}(\mathrm{R}_f, \mathrm{R}_g) := \inf_{\phi,\psi} \max \left\{ \frac{1}{2} D(\phi, \psi),\ \|f - g \circ \phi\|_\infty,\ \|f \circ \psi - g\|_\infty \right\}, \tag{4}$$

where:

- $\phi : \mathrm{R}_f \to \mathrm{R}_g$ *and* $\psi : \mathrm{R}_g \to \mathrm{R}_f$ *are continuous maps,*
- $D(\phi, \psi) = \sup \{|d_f(x, x') - d_g(y, y')| \text{ such that } (x, y), (x', y') \in C(\phi, \psi)\}$, *where:*
  - $C(\phi, \psi) = \{(x, \phi(x)) \mid x \in \mathrm{R}_f\} \cup \{(\psi(y), y) \mid y \in \mathrm{R}_g\}$,

– $d_f(x,x') = \min_{\pi: x \to x'} \left\{ \max_{t\in[0,1]} f\circ\pi(t) - \min_{t\in[0,1]} f\circ\pi(t) \right\}$, *where* $\pi : [0,1] \to \mathrm{R}_f$ *is a continuous path from* x *to* x' *in* R_f *(*$\pi(0) = x$ *and* $\pi(1) = x'$*),*

– $d_g(y,y') = \min_{\pi: y \to y'} \left\{ \max_{t\in[0,1]} g\circ\pi(t) - \min_{t\in[0,1]} g\circ\pi(t) \right\}$, *where* $\pi : [0,1] \to \mathrm{R}_g$ *is a continuous path from* y *to* y' *in* R_g *(*$\pi(0) = y$ *and* $\pi(1) = y'$*).*

Bauer et al. [4] related these distances as follows:

Theorem 2.7. *The following inequality holds:* $d_{\mathrm{B}}(\mathrm{R}_f, \mathrm{R}_g) \leq 3\, d_{\mathrm{FD}}(\mathrm{R}_f, \mathrm{R}_g)$.

This result can be improved using the end of Section 3.4 of [7], then noting that level set diagrams and extended diagrams are essentially the same [10], and finally Lemma 9 of [5]:

Theorem 2.8. *The following inequality holds:* $d_{\mathrm{B}}(\mathrm{R}_f, \mathrm{R}_g) \leq 2\, d_{\mathrm{FD}}(\mathrm{R}_f, \mathrm{R}_g)$.

We emphasize that, even though Theorem 2.8 allows us to improve on the constants of our main result—see Theorem 3.1, the reduction from $3\, d_{\mathrm{FD}}(\mathrm{R}_f, \mathrm{R}_g)$ in Theorem 2.7 to $2\, d_{\mathrm{FD}}(\mathrm{R}_f, \mathrm{R}_g)$ in Theorem 2.8 is not fundamental for our analysis and proofs.

Since the bottleneck distance is only a pseudo-metric—see Figure 1, the inequality given by Theorem 2.8 cannot be turned into an equivalence result. However, for any pair of Reeb graphs $\mathrm{R}_f, \mathrm{R}_g$ that have the same extended persistence diagram $\mathrm{Dg}(f) = \mathrm{Dg}(g)$, and that are at positive functional distortion distance from each other, every continuous path in d_{FD} from R_f to R_g will perturb the points of $\mathrm{Dg}(f)$ and eventually drive them back to their initial position, suggesting first that d_{B} is locally equivalent to d_{FD}—see Theorem 3.1 in Section 3, but also that, even though $d_{\mathrm{B}}(\mathrm{R}_f, \mathrm{R}_g) = 0$, the intrinsic metric $\hat{d}_{\mathrm{B}}(\mathrm{R}_f, \mathrm{R}_g)$ induced by d_{B} is positive—see Theorem 4.2 in Section 4.

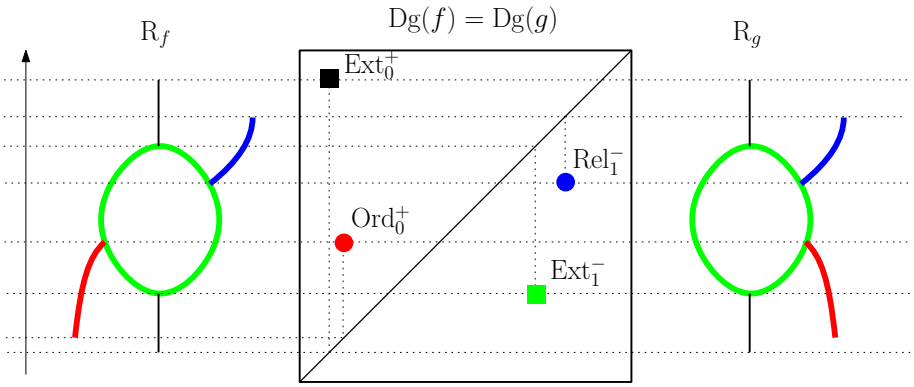

Figure 1: Example of two different Reeb graphs R_f and R_g that have the same extended persistence diagram $\mathrm{Dg}(f) = \mathrm{Dg}(g)$. These graphs are at bottleneck distance 0 from each other, while their functional distortion distance is positive.

3 Local Equivalence

Let $a_f = \min_{1\leq i\leq n} a_{i+1} - a_i > 0$ and $a_g = \min_{1\leq j\leq m} b_{j+1} - b_j > 0$. In this section, we show the following local equivalence theorem:

Theorem 3.1. *Let $K \in (0, 1/22]$. If $d_{\mathrm{FD}}(\mathrm{R}_f, \mathrm{R}_g) \leq \max\{a_f, a_g\}/(8(1+22K))$, then:*

$$K d_{\mathrm{FD}}(\mathrm{R}_f, \mathrm{R}_g) \leq d_{\mathrm{B}}(\mathrm{R}_f, \mathrm{R}_g) \leq 2 d_{\mathrm{FD}}(\mathrm{R}_f, \mathrm{R}_g).$$

Note that the notion of locality used here is slightly different from the usual one. On the one hand, the equivalence does not hold for any arbitrary pair of Reeb graphs inside a neighborhood of some fixed Reeb graph, but rather for any pair involving the fixed graph. On the other hand, the constants in the equivalence are independent of the pair of Reeb graphs considered. The upper bound on $d_{\mathrm{B}}(\mathrm{R}_f, \mathrm{R}_g)$ is given by Theorem 2.8 and always holds. The aim of this section is to prove the lower bound.

Convention: We assume w.l.o.g. that $\max\{a_f, a_g\} = a_f$, and we let $\varepsilon = d_{\mathrm{FD}}(\mathrm{R}_f, \mathrm{R}_g)$.

3.1 Proof of Theorem 3.1

Let $K \in (0, 1/22]$. The proof proceeds by contradiction. Assuming $d_{\mathrm{B}}(\mathrm{R}_f, \mathrm{R}_g) < K\varepsilon$, where $\varepsilon = d_{\mathrm{FD}}(\mathrm{R}_f, \mathrm{R}_g) < a_f/(8(1+22K))$, we progressively transform R_g into some other Reeb graph $\mathrm{R}_{g'}$ (Definition 3.4) that satisfies both $d_{\mathrm{FD}}(\mathrm{R}_g, \mathrm{R}_{g'}) < \varepsilon$ (Proposition 3.6) and $d_{\mathrm{FD}}(\mathrm{R}_f, \mathrm{R}_{g'}) = 0$ (Proposition 3.7). The contradiction follows from the triangle inequality.

3.1.1 Graph Transformation

The graph transformation is defined as the composition of the *simplification operator* from [4] and the *Merge operator*[1] from [11]. We refer the reader to these articles for the precise definitions. Below we merely recall their main properties. Given a set $S \subseteq X$ and a scalar $\alpha > 0$, we recall that $S^\alpha = \{x \in X \mid d(x, S) \leq \alpha\}$ denotes the α-offset of S.

Lemma 3.2 (Theorem 7.3 and following remark in [4])**.** *Given $\alpha > 0$, the simplification operator $S_\alpha : \mathsf{Reeb} \to \mathsf{Reeb}$ takes any Reeb graph R_h to $\mathrm{R}_{h'} = S_\alpha(\mathrm{R}_h)$ such that $\mathrm{Dg}(h') \cap \Delta^{\alpha/2} = \emptyset$ and $d_{\mathrm{B}}(\mathrm{R}_h, \mathrm{R}_{h'}) \leq 2\, d_{\mathrm{FD}}(\mathrm{R}_h, \mathrm{R}_{h'}) \leq 4\alpha$.*

Lemma 3.3 (Theorem 2.5 and Lemma 4.3 in [11])**.** *Given $a \leq b$, the merge operator $\mathrm{Merge}_{a,b} : \mathsf{Reeb} \to \mathsf{Reeb}$ takes any Reeb graph R_h to $\mathrm{R}_{h'} = \mathrm{Merge}_{a,b}(\mathrm{R}_h)$ such that $\mathrm{Dg}(h')$ is obtained from $\mathrm{Dg}(h)$ through the following snapping principle (see Figure 2 for an illustration):*

$$(x, y) \in \mathrm{Dg}(h) \mapsto (x', y') \in \mathrm{Dg}(h') \text{ where } x' = \begin{cases} x \text{ if } x \notin [a, b] \\ \frac{a+b}{2} \text{ otherwise} \end{cases} \quad \text{and similarly for } y'.$$

Definition 3.4. *Let R_f be a fixed Reeb graph with critical values $\{a_1, \cdots, a_n\}$. Given $\alpha > 0$, the full transformation $F_\alpha : \mathsf{Reeb} \to \mathsf{Reeb}$ is defined as $F_\alpha = \mathrm{Merge}_{9\alpha} \circ S_{2\alpha}$, where $\mathrm{Merge}_{9\alpha} = \mathrm{Merge}_{a_n - 9\alpha,\, a_n + 9\alpha} \circ \ldots \circ \mathrm{Merge}_{a_1 - 9\alpha,\, a_1 + 9\alpha}$. See Figure 3 for an illustration.*

[1] Strictly speaking, the output of our Merge is the Reeb graph of the output of the Merge from [11].

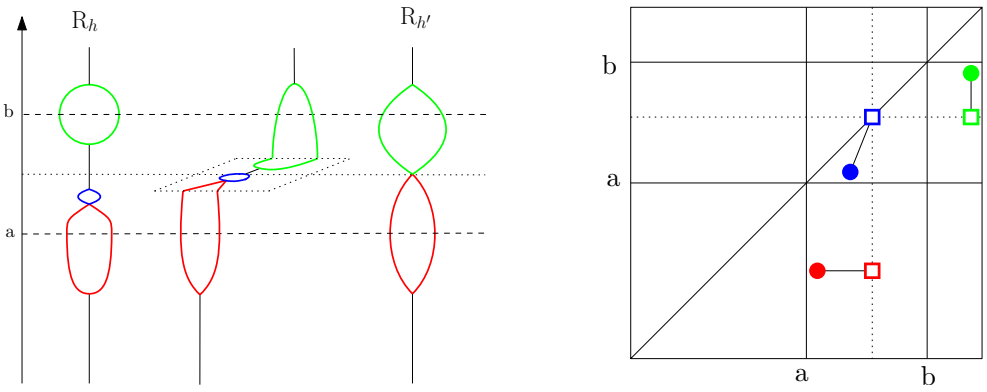

Figure 2: Left: effect of $\mathrm{Merge}_{a,b}$ on a Reeb graph R_h. Right: Effect on its persistence diagram.

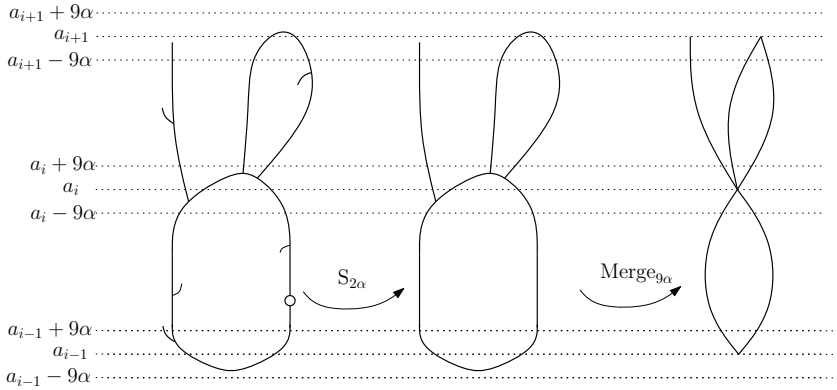

Figure 3: Illustration of F_α.

3.1.2 Properties of the transformed graph

Let $\mathrm{R}_f, \mathrm{R}_g \in \mathsf{Reeb}$ such that $d_{\mathrm{B}}(\mathrm{R}_f, \mathrm{R}_g) < K\varepsilon$ where $\varepsilon = d_{\mathrm{FD}}(\mathrm{R}_f, \mathrm{R}_g) < a_f/(8(1+22K))$. Letting $\mathrm{R}_{g'} = F_{K\varepsilon}(\mathrm{R}_g)$, we want to show both that $d_{\mathrm{FD}}(\mathrm{R}_g, \mathrm{R}_{g'}) < 22K\varepsilon < \varepsilon$ and $d_{\mathrm{FD}}(\mathrm{R}_f, \mathrm{R}_{g'}) = 0$, which will lead to a contradiction as mentioned previously.

Let $B_\infty(\cdot,\cdot)$ denote balls in the ℓ_∞-norm.

Lemma 3.5. *Let* $\mathrm{R}_h = S_{2K\varepsilon}(\mathrm{R}_g)$. *Under the above assumptions, one has*

$$\mathrm{Dg}(h) \subseteq \bigcup\nolimits_{\tau \in \mathrm{Dg}(f)} B_\infty(\tau, 9K\varepsilon). \tag{5}$$

Proof. Since $d_{\mathrm{B}}(\mathrm{R}_f, \mathrm{R}_g) < K\varepsilon$, we have $\mathrm{Dg}(g) \subseteq \bigcup_{\tau\in\mathrm{Dg}(f)} B_\infty(\tau, K\varepsilon) \cup \Delta^{K\varepsilon}$. Since $\mathrm{R}_h = S_{2K\varepsilon}(\mathrm{R}_g)$, it follows from Lemma 3.2 that $d_{\mathrm{B}}(\mathrm{Dg}(h), \mathrm{Dg}(g)) \leq 8K\varepsilon$. Moreover, since every persistence pair in $\mathrm{Dg}(g) \cap \Delta^{K\varepsilon}$ is removed by $S_{2K\varepsilon}$, it results that:

$\mathrm{Dg}(h) \subseteq \bigcup_{\tau\in\mathrm{Dg}(g)\setminus\Delta^{K\varepsilon}} B_\infty(\tau, 8K\varepsilon) \subseteq \bigcup_{\tau\in\mathrm{Dg}(f)} B_\infty(\tau, 9K\varepsilon)$. □

Now we bound $d_{\mathrm{FD}}(\mathrm{R}_{g'}, \mathrm{R}_g)$. Recall that, given an arbitrary Reeb graph R_h, with critical values $\mathrm{Crit}(h) = \{c_1, ..., c_p\}$, if C is a cc of $h^{-1}(I)$, where I is an open interval such that $\exists c_i, c_{i+1}$ s.t. $I \subseteq (c_i, c_{i+1})$, then C is a *topological arc*, i.e. homeomorphic to an open interval.

Proposition 3.6. *Under the same assumptions as above, one has* $d_{\mathrm{FD}}(\mathrm{R}_g, \mathrm{R}_{g'}) < 22K\epsilon$.

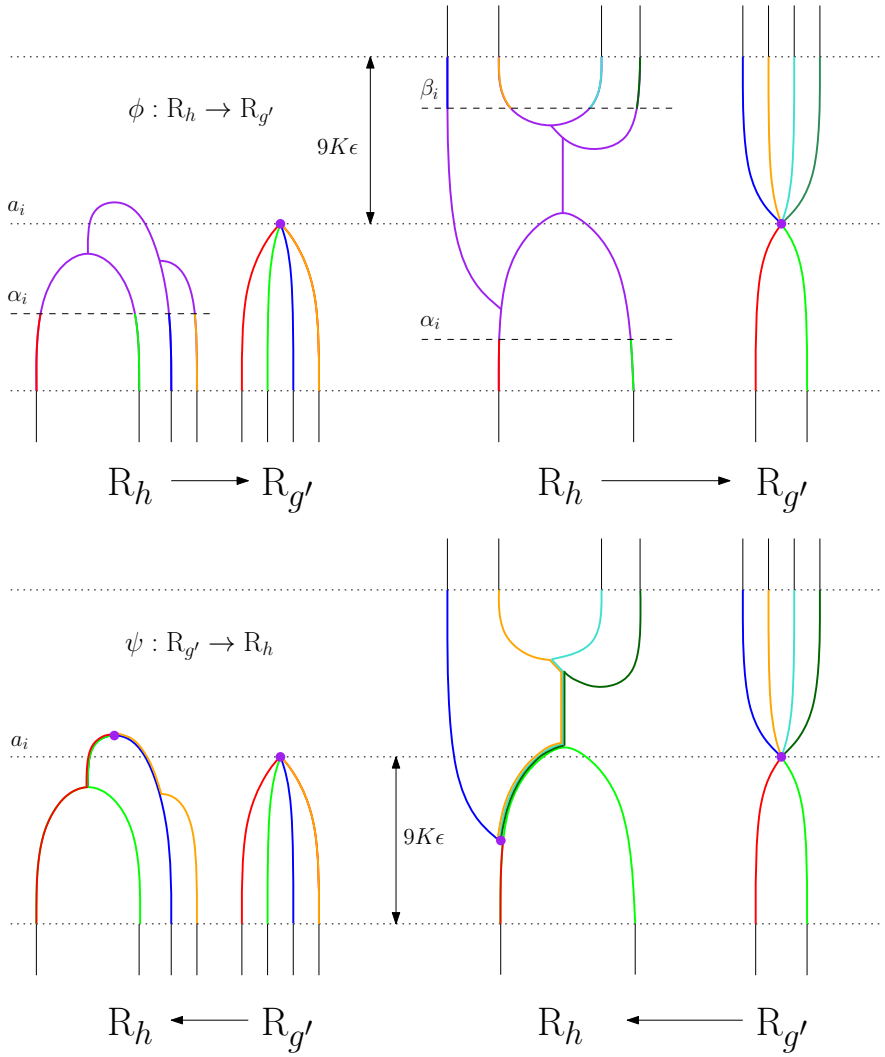

Figure 4: The effects of ϕ and ψ around a specific critical value a_i of f. Segments are matched according to their colors (up to reparameterization).

Proof. Let $\mathrm{R}_h = S_{2K\varepsilon}(\mathrm{R}_g)$. We have $d_{\mathrm{FD}}(\mathrm{R}_{g'}, \mathrm{R}_g) \leq d_{\mathrm{FD}}(\mathrm{R}_{g'}, \mathrm{R}_h) + d_{\mathrm{FD}}(\mathrm{R}_h, \mathrm{R}_g)$ by the triangle inequality. It suffices therefore to bound both $d_{\mathrm{FD}}(\mathrm{R}_{g'}, \mathrm{R}_h)$ and $d_{\mathrm{FD}}(\mathrm{R}_h, \mathrm{R}_g)$. By Lemma 3.2, we have $d_{\mathrm{FD}}(\mathrm{R}_h, \mathrm{R}_g) < 4K\varepsilon$. Now, recall from (5) that the points of the extended persistence diagram of R_h are included in $\bigcup_{\tau \in \mathrm{Dg}(f)} B_\infty(\tau, 9K\varepsilon)$. Moreover, since $\mathrm{R}_{g'} = \mathrm{Merge}_{9K\varepsilon}(\mathrm{R}_h)$, $\mathrm{R}_{g'}$ and R_h are composed of the same number of arcs in each $[a_i + 9K\varepsilon, a_{i+1} - 9K\varepsilon]$. Hence, we can define explicit continuous maps $\phi : \mathrm{R}_h \to \mathrm{R}_{g'}$ and $\psi : \mathrm{R}_{g'} \to \mathrm{R}_h$ as depicted in Figure 4. More precisely, since R_h and $\mathrm{R}_{g'}$ are composed of the same number of arcs in each $[a_i + 9K\varepsilon, a_{i+1} - 9K\varepsilon]$, we only need to specify ϕ and ψ inside each interval $(a_i - 9K\varepsilon, a_i + 9K\varepsilon)$ and then ensure that the piecewise-defined maps are assembled consistently. Since the critical values of R_h are within distance less that $9K\varepsilon$ of the critical values of f, there exist two levels $a_i - 9K\varepsilon < \alpha_i \leq \beta_i < a_i + 9K\varepsilon$ such that R_h is only

composed of arcs in $(a_i - 9K\varepsilon, \alpha_i]$ and $[\beta_i, a_i + 9K\varepsilon)$ for each i (dashed lines in Figure 4). For any cc C of $h^{-1}((a_i - 9K\varepsilon, a_i + 9K\varepsilon))$, the map ϕ sends all points of $C \cap h^{-1}([\alpha_i, \beta_i])$ to the corresponding critical point y_C created by the Merge in $\mathrm{R}_{g'}$, and it extends the arcs of $C \cap h^{-1}((a_i - 9K\varepsilon, \alpha_i])$ (resp. $C \cap h^{-1}([\beta_i, a_i + 9K\varepsilon))$) into arcs of $(g')^{-1}([a_i - 9K\varepsilon, a_i])$ (resp. $(g')^{-1}([a_i, a_i + 9K\varepsilon])$). In return, the map ψ sends the critical point y_C to an arbitrary point of C. Then, since the Merge operation preserves connected components, for each arc A' of $(g')^{-1}((a_i - 9K\varepsilon, a_i + 9K\varepsilon))$ connected to y_C, there is at least one corresponding path A in R_h whose endpoint in $h^{-1}(a_i - 9K\varepsilon)$ or $h^{-1}(a_i + 9K\varepsilon)$ matches with the one of A' (see the colors in Figure 4). Hence ψ sends A' to A and the piecewise-defined maps are assembled consistently.
Let us bound the three terms in the $\max\{\cdots\}$ in (4) with this choice of maps ϕ, ψ:

- We first bound $\|g' - h \circ \psi\|_\infty$. Let $x \in \mathrm{R}_{g'}$. Either $g'(x) \in \bigcup_{i \in \{1,\dots,n-1\}} [a_i + 9K\varepsilon, a_{i+1} - 9K\varepsilon]$, and in this case we have $g'(x) = h(\psi(x))$ by definition of ψ; or, there is $i_0 \in \{1, ..., n\}$ such that $g'(x) \in (a_{i_0} - 9K\varepsilon, a_{i_0} + 9K\varepsilon)$ and then $h(\psi(x)) \in (a_{i_0} - 9K\varepsilon, a_{i_0} + 9K\varepsilon)$. In both cases $|g'(x) - h \circ \psi(x)| < 18K\varepsilon$. Hence, $\|g' - h \circ \psi\|_\infty < 18K\varepsilon$.

- Since the previous proof is symmetric in h and g', one also has $\|h - g' \circ \phi\|_\infty < 18K\varepsilon$.

- We now bound $D(\phi, \psi)$. Let $(x, \phi(x)), (\psi(y), y) \in C(\phi, \psi)$ (the cases $(x, \phi(x)), (x', \phi(x'))$ and $(\psi(y), y), (\psi(y'), y')$ are similar). Let $\pi_{g'} : [0, 1] \to \mathrm{R}_{g'}$ be a continuous path from $\phi(x)$ to y which achieves $d_{g'}(\phi(x), y)$.

  - Assume $h(x) \in \bigcup_{i \in \{1,\dots,n-1\}} [a_i + 9K\varepsilon, a_{i+1} - 9K\varepsilon]$. Then one has $\psi \circ \phi(x) = x$. Hence, $\pi_h := \psi \circ \pi_{g'}$ is a valid path from x to $\psi(y)$. Moreover, since $\|g' - h \circ \psi\|_\infty < 18K\varepsilon$, it follows that
  $$\begin{aligned} &\max\ \mathrm{im}(h \circ \pi_h) < \max\ \mathrm{im}(g' \circ \pi_{g'}) + 18K\varepsilon, \\ &\min\ \mathrm{im}(h \circ \pi_h) > \min\ \mathrm{im}(g' \circ \pi_{g'}) - 18K\varepsilon. \end{aligned} \tag{6}$$
  Hence, one has
  $$\begin{aligned} &d_h(x, \psi(y)) \leq \max\ \mathrm{im}(h \circ \pi_h) - \min\ \mathrm{im}(h \circ \pi_h) < d_{g'}(\phi(x), y) + 36K\varepsilon, \\ &-d_h(x, \psi(y)) \geq \min\ \mathrm{im}(h \circ \pi_h) - \max\ \mathrm{im}(h \circ \pi_h) > -d_{g'}(\phi(x), y) - 36K\varepsilon. \end{aligned}$$
  This shows that $|d_h(x, \psi(y)) - d_{g'}(\phi(x), y)| < 36K\varepsilon$.

  - Assume that there is $i_0 \in \{1, ..., n\}$ such that $h(x) \in (a_{i_0} - 9K\varepsilon, a_{i_0} + 9K\varepsilon)$. Then, by definition of ϕ, ψ, we have $g'(\phi(x)) \in (a_{i_0} - 9K\varepsilon, a_{i_0} + 9K\varepsilon)$, and, since ϕ and ψ preserve connected components, there is a path $\pi'_h : [0, 1] \to \mathrm{R}_h$ from x to $\psi \circ \phi(x)$ within the interval $(a_{i_0} - 9K\varepsilon, a_{i_0} + 9K\varepsilon)$, which itself is included in the interior of the offset $\mathrm{im}(g' \circ \pi_{g'})^{18K\varepsilon}$. Let now π_h be the concatenation of π'_h with $\psi \circ \pi_{g'}$, which goes from x to $\psi(y)$. Since $\|g' - h \circ \psi\| < 18K\varepsilon$, it follows that $\mathrm{im}(h \circ \psi \circ \pi_{g'}) \subseteq \mathrm{int}\ \mathrm{im}(g' \circ \pi_{g'})^{18K\varepsilon}$, and since $\mathrm{im}(h \circ \pi_h) = \mathrm{im}(h \circ \pi'_h) \cup \mathrm{im}(h \circ \psi \circ \pi_{g'})$ by concatenation, one finally has $\mathrm{im}(h \circ \pi_h) \subseteq \mathrm{int}\ \mathrm{im}(g' \circ \pi_{g'})^{18K\varepsilon}$. Hence, the inequalities of (6) hold, implying that $|d_h(x, \psi(y)) - d_{g'}(\phi(x), y)| < 36K\varepsilon$.

  Since these inequalities hold for any couples $(x, \phi(x))$ and $(\psi(y), y)$, we deduce that $D(\phi, \psi) \leq 36K\varepsilon$.

Thus, $d_{\mathrm{FD}}(\mathrm{R}_h, \mathrm{R}_g) < 4K\varepsilon$ and $d_{\mathrm{FD}}(\mathrm{R}_h, \mathrm{R}_{g'}) \leq 18K\varepsilon$, so $d_{\mathrm{FD}}(\mathrm{R}_{g'}, \mathrm{R}_g) < 22K\varepsilon$ as desired. □

Now we show that $\mathrm{R}_{g'}$ is isomorphic to R_f (i.e. it lies at functional distortion distance 0).

Proposition 3.7. *Under the same assumptions as above, one has* $d_{\mathrm{FD}}(\mathrm{R}_f, \mathrm{R}_{g'}) = 0$.

Proof. First, recall from (5) that the points of the extended persistence diagram of R_h are included in $\bigcup_{\tau \in \mathrm{Dg}(f)} B_\infty(\tau, 9K\varepsilon)$. Since $\mathrm{R}_{g'} = \mathrm{Merge}_{9K\varepsilon}(\mathrm{R}_h)$, it follows from Lemma 3.3 that $\mathrm{Crit}(g') \subseteq \mathrm{Crit}(f)$. Hence, both $\mathrm{R}_{g'}$ and R_f are composed of arcs in each (a_i, a_{i+1}).

Now, we show that, for each i, the number of arcs of $(g')^{-1}((a_i, a_{i+1}))$ and $f^{-1}((a_i, a_{i+1}))$ are the same. By the triangle inequality and Proposition 3.6, we have:

$$d_{\mathrm{FD}}(\mathrm{R}_f, \mathrm{R}_{g'}) \leq d_{\mathrm{FD}}(\mathrm{R}_f, \mathrm{R}_g) + d_{\mathrm{FD}}(\mathrm{R}_g, \mathrm{R}_{g'}) < (1 + 22K)\varepsilon. \tag{7}$$

Let $\phi : \mathrm{R}_f \to \mathrm{R}_{g'}$ and $\psi : \mathrm{R}_{g'} \to \mathrm{R}_f$ be optimal continuous maps that achieve $d_{\mathrm{FD}}(\mathrm{R}_f, \mathrm{R}_{g'})$. Let $i \in \{1, ..., n-1\}$. Assume that there are more arcs of $f^{-1}((a_i, a_{i+1}))$ than arcs of $(g')^{-1}((a_i, a_{i+1}))$. For every arc A of $f^{-1}((a_i, a_{i+1}))$, let $x_A \in A$ such that $f(x_A) = \bar{a} = \frac{1}{2}(a_i + a_{i+1})$. First, note that $\phi(x_A)$ must belong to an arc of $(g')^{-1}((a_i, a_{i+1}))$. Indeed, since $\|f - g' \circ \phi\|_\infty < (1 + 22K\varepsilon)$, one has $g'(\phi(x_A)) \in (\bar{a} - (1 + 22K)\varepsilon, \bar{a} + (1 + 22K)\varepsilon) \subseteq (a_i, a_{i+1})$. Then, according to the pigeonhole principle, there exist $x_A, x_{A'}$ such that $\phi(x_A)$ and $\phi(x_{A'})$ belong to the same arc of $(g')^{-1}((a_i, a_{i+1}))$.

- Since x_A and $x_{A'}$ do not belong to the same arc, we have $d_f(x_A, x_{A'}) > a_f/2$.
- Now, since $\|f - g' \circ \phi\|_\infty < (1 + 22K)\varepsilon$ and $\phi(x_A), \phi(x_{A'})$ belong to the same arc of $(g')^{-1}((a_i, a_{i+1}))$, we also have $d_{g'}(\phi(x_A), \phi(x_{A'})) < 2(1 + 22K)\varepsilon$ (see Figure 5).

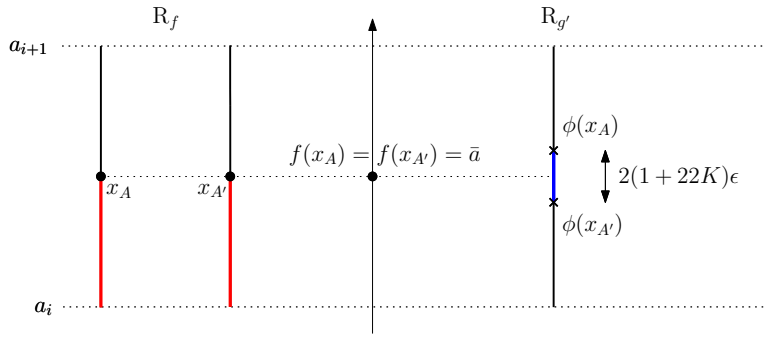

Figure 5: Any path between x_A and $x_{A'}$ must contain the red segments, and the blue segment is a particular path between $\phi(x_A)$ and $\phi(x_{A'})$.

Hence, $D(\phi, \psi) \geq |d_f(x_A, x_{A'}) - d_{g'}(\phi(x_A), \phi(x_{A'}))| > a_f/2 - 2(1 + 22K)\varepsilon$, which is greater than $2(1 + 22K)\varepsilon$ because $\varepsilon < a_f/(8(1 + 22K))$. Thus, $d_{\mathrm{FD}}(\mathrm{R}_f, \mathrm{R}_{g'}) > (1 + 22K)\varepsilon$, which leads to a contradiction with (7). This means that there cannot be more arcs in $f^{-1}((a_i, a_{i+1}))$ than in $(g')^{-1}((a_i, a_{i+1}))$. Since the proof is symmetric in f and g', the numbers of arcs in $(g')^{-1}((a_i, a_{i+1}))$ and in $f^{-1}((a_i, a_{i+1}))$ are actually the same.

Finally, we show that the attaching maps of these arcs are also the same. In this particular graph setting, this is equivalent to showing that corresponding arcs in R_f and $\mathrm{R}_{g'}$ have the same endpoints. Let a_i be a critical value. Let $A^-_{f,i}$ and $A^+_{f,i}$ (resp. $A^-_{g',i}$ and $A^+_{g',i}$) be the sets of arcs in $f^{-1}((a_{i-1}, a_i))$ and $f^{-1}((a_i, a_{i+1}))$ (resp. $(g')^{-1}((a_{i-1}, a_i))$ and $(g')^{-1}((a_i, a_{i+1}))$). Morevover, we let ζ^i_f and ξ^i_f (resp. $\zeta^i_{g'}$ and $\xi^i_{g'}$) be the corresponding attaching maps that send arcs to their

endpoints in $f^{-1}(a_i)$ (resp. $(g')^{-1}(a_i)$). Let $A, B \in A^-_{f,i}$. We define an equivalence relation $\sim_{f,i}$ between A and B by: $A \sim_{f,i} B$ iff $\zeta^i_f(A) = \zeta^i_f(B)$, i.e. the endpoints of the arcs in the critical slice $f^{-1}(a_i)$ are the same. Similarly, $C, D \in A^+_{f,i}$ are equivalent if and only if $\xi^i_f(C) = \xi^i_f(D)$. One can define $\sim_{g',i}$ in the same way. To show that the attaching maps of R_f and $\mathrm{R}_{g'}$ are the same, we need to find a bijection b between the arcs of R_f and $\mathrm{R}_{g'}$ such that $A \sim_{f,i} B \Leftrightarrow b(A) \sim_{g',i} b(B)$ for each i.

We will now define b then check that it satisfies the condition. Recall from (7) that $d_{\mathrm{FD}}(\mathrm{R}_f, \mathrm{R}_{g'}) < (1+22K)\varepsilon$. Hence there exists a continuous map $\phi : \mathrm{R}_f \to \mathrm{R}_{g'}$ such that $\|f - g' \circ \phi\|_\infty < (1+22K)\varepsilon$. This map induces a bijection b between the arcs of R_f and $\mathrm{R}_{g'}$. Indeed, given an arc $A \in A^-_{f,i}$, let $x \in A$ such that $f(x) = \bar{a} = \frac{1}{2}(a_{i-1} + a_i)$. We define $b(A)$ as the arc of $A^-_{g,i}$ that contains $\phi(x)$. The map b is well-defined since $g' \circ \phi(x) \in [\bar{a} - (1+22K)\varepsilon, \bar{a} + (1+22K)\varepsilon] \subseteq (a_{i-1}, a_i)$, hence $\phi(x)$ must belong to an arc of $(g')^{-1}((a_{i-1}, a_i))$. Let us show that $b(A) \sim_{g',i} b(B) \Rightarrow A \sim_{f,i} B$. Assume there exist $A, B \in A^-_{f,i}$ (the treatment of $A, B \in A^+_{f,i}$ is similar) such that $A \not\sim_{f,i} B$ and $b(A) \sim_{g',i} b(B)$. Let $x = \zeta^i_f(A)$ and $y = \zeta^i_f(B)$. Then we have $d_f(x, y) \geq a_f$ while $d_{g'}(\phi(x), \phi(y)) < 2(1 + 22K)\varepsilon$ (see Figure 6). Hence $|d_f(x, y) - d_{g'}(\phi(x), \phi(y))| > a_f - 2(1 + 22K)\varepsilon > 2(1 + 22K)\varepsilon$, so $d_{\mathrm{FD}}(\mathrm{R}_f, \mathrm{R}_{g'}) > (1+22K)\varepsilon$, which leads to a contradiction with (7). The same argument applies to show that $A \sim_{f,i} B \Rightarrow b(A) \sim_{g',i} b(B)$. □

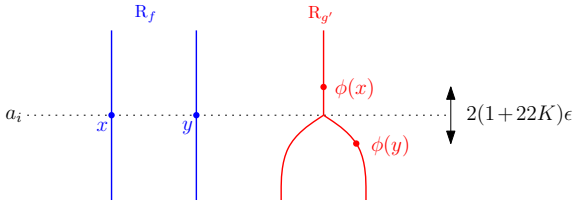

Figure 6: Any path from x to y must go through an entire arc, hence $d_f(x, y) \geq a_f$. On the contrary, there exists a direct path between $\phi(x)$ and $\phi(y)$, hence $d_{g'}(\phi(x), \phi(y)) < 2(1 + 22K)\varepsilon$.

4 Induced Intrinsic Metrics

In this section we leverage the local equivalence given by Theorem 3.1 to derive a global equivalence between the intrinsic metrics $\hat{d}_{\mathrm{B}}$ and $\hat{d}_{\mathrm{FD}}$ induced by d_{B} and d_{FD}. Note that we already know $\hat{d}_{\mathrm{FD}}$ to be equivalent to $\hat{d}_{\mathrm{GH}}$ and $\hat{d}_{\mathrm{I}}$ since d_{FD} is equivalent to d_{GH} and d_{I}. To the best of our knowledge, the question whether d_{FD}, d_{I} or d_{GH} is intrinsic on the space of Reeb graphs has not been settled, although d_{GH} itself is known to be intrinsic on the larger space of compact metric spaces—see e.g. [18].

Convention. In the following, whatever the metric $d : \mathsf{Reeb} \times \mathsf{Reeb} \to \mathbb{R}_+$ under consideration, we define the class of *admissible paths* in Reeb to be those maps $\gamma : [0, 1] \to \mathsf{Reeb}$ that are continuous in d_{FD}. This makes sense when d is either d_{FD} itself, d_{GH}, or d_{I}, all of which are equivalent to d_{FD} and therefore have the same continuous maps $\gamma : [0, 1] \to \mathsf{Reeb}$. In the case $d = d_{\mathrm{B}}$ our convention means restricting the class of admissible paths to a strict subset of the maps $\gamma : [0, 1] \to \mathsf{Reeb}$ that are continuous in d (by Theorem 2.8), which is required by some of our following claims.

Definition 4.1. *Let* $d : \mathsf{Reeb} \times \mathsf{Reeb} \to \mathbb{R}_+$ *be a metric on* Reeb*. Let* $\mathrm{R}_f, \mathrm{R}_g \in \mathsf{Reeb}$*, and* $\gamma : [0,1] \to \mathsf{Reeb}$ *be an admissible path such that* $\gamma(0) = \mathrm{R}_f$ *and* $\gamma(1) = \mathrm{R}_g$*. The* length *of* γ *induced by* d *is defined as* $L_d(\gamma) = \sup_{n,\Sigma} \sum_{i=0}^{n-1} d(\gamma(t_i), \gamma(t_{i+1}))$ *where* n *ranges over* $\mathbb{N}$ *and* Σ *ranges over all partitions* $0 = t_0 \leq t_1 \leq ... \leq t_n = 1$ *of* $[0,1]$*. The* intrinsic metric induced by d*, denoted* $\hat{d}$*, is defined by* $\hat{d}(\mathrm{R}_f, \mathrm{R}_g) = \inf_\gamma \mathrm{L_d}(\gamma)$ *where* γ *ranges over all admissible paths* $\gamma : [0,1] \to \mathsf{Reeb}$ *such that* $\gamma(0) = \mathrm{R}_f$ *and* $\gamma(1) = \mathrm{R}_g$*.*

The following result is, in our view, the starting point for the study of intrinsic metrics over the space of Reeb graphs. It comes as a consequence of the (local or global) equivalences between d_{B} and d_{FD} stated in Theorems 2.8 and 3.1. The intuition is that integrating two locally equivalent metrics along the same path using sufficiently small integration steps yields the same total length up to a constant factor, hence the global equivalence between the induced intrinsic metrics[2].

Theorem 4.2. $\hat{d}_{\mathrm{B}}$ *and* $\hat{d}_{\mathrm{FD}}$ *are globally equivalent. Specifically, for any* $\mathrm{R}_f, \mathrm{R}_g \in \mathsf{Reeb}$*,*

$$\hat{d}_{\mathrm{FD}}(\mathrm{R}_f, \mathrm{R}_g)/22 \leq \hat{d}_{\mathrm{B}}(\mathrm{R}_f, \mathrm{R}_g) \leq 2\, \hat{d}_{\mathrm{FD}}(\mathrm{R}_f, \mathrm{R}_g). \tag{8}$$

Proof. We first show that $\hat{d}_{\mathrm{B}}(\mathrm{R}_f, \mathrm{R}_g) \leq 2\, \hat{d}_{\mathrm{FD}}(\mathrm{R}_f, \mathrm{R}_g)$. Let γ be an admissible path and let $\Sigma = \{t_0, ..., t_n\}$ be a partition of $[0,1]$. Then, by Theorem 2.8,

$$\sum_{i=0}^{n-1} d_{\mathrm{FD}}(\gamma(t_i), \gamma(t_{i+1})) \geq \frac{1}{2} \sum_{i=0}^{n-1} d_{\mathrm{B}}(\gamma(t_i), \gamma(t_{i+1})).$$

Since this is true for any partition Σ of any finite size n, it follows that

$$L_{d_{\mathrm{FD}}}(\gamma) \geq \frac{1}{2} L_{d_{\mathrm{B}}}(\gamma) \geq \frac{1}{2} \hat{d}_{\mathrm{B}}(\mathrm{R}_f, \mathrm{R}_g).$$

Again, this inequality holds for any admissible path γ, so $\hat{d}_{\mathrm{B}}(\mathrm{R}_f, \mathrm{R}_g) \leq 2\hat{d}_{\mathrm{FD}}(\mathrm{R}_f, \mathrm{R}_g)$.
We now show that $\hat{d}_{\mathrm{FD}}(\mathrm{R}_f, \mathrm{R}_g)/22 \leq \hat{d}_{\mathrm{B}}(\mathrm{R}_f, \mathrm{R}_g)$. Let γ be an admissible path and $\Sigma = \{t_0, ..., t_n\}$ a partition of $[0,1]$. We claim that there is a refinement of Σ (i.e. a partition $\Sigma' = \{t'_0, ..., t'_m\} \supseteq \Sigma$ for some $m \geq n$) such that $d_{\mathrm{FD}}(\gamma(t'_j), \gamma(t'_{j+1})) < \max\{c_{t'_j}, c_{t'_{j+1}}\}/16$ for all $j \in \{0, ..., m-1\}$, where $c_t > 0$ denotes the minimal distance between consecutive critical values of $\gamma(t)$. Indeed, since γ is continuous in d_{FD}, for any $t \in [0,1]$ there exists $\delta_t > 0$ such that $d_{\mathrm{FD}}(\gamma(t), \gamma(t')) < c_t/16$ for all $t' \in [0,1]$ with $|t - t'| < \delta_t$. Consider the open cover $\{(\max\{0, t - \delta_t/2\}, \min\{1, t + \delta_t/2\})\}_{t \in [0,1]}$ of $[0,1]$. Since $[0,1]$ is compact, there exists a finite subcover containing all the intervals $(t_i - \delta_{t_i}/2, t_i + \delta_{t_i}/2)$ for $t_i \in \Sigma$. Assume w.l.o.g. that this subcover is minimal (if it is not, then reduce the δ_{t_i} as much as needed). Let then $\Sigma' = \{t'_0, ..., t'_m\} \supseteq \Sigma$ be the partition of $[0,1]$ given by the midpoints of the intervals in this subcover, sorted by increasing order. Since the subcover is minimal, we have $t'_{j+1} - t'_j < (\delta_{t'_j} + \delta_{t'_{j+1}})/2 < \max\{\delta_{t'_j}, \delta_{t'_{j+1}}\}$ hence $d_{\mathrm{FD}}(\gamma(t'_j), \gamma(t'_{j+1})) < \max\{c_{t'_j}, c_{t'_{j+1}}\}/16$ for each $j \in \{0, m-1\}$. It follows that

$$\begin{aligned}
\sum_{i=0}^{n-1} d_{\mathrm{FD}}(\gamma(t_i), \gamma(t_{i+1})) &\leq \sum_{j=0}^{m-1} d_{\mathrm{FD}}(\gamma(t'_j), \gamma(t'_{j+1})) \text{ by the triangle inequality since } \Sigma' \supseteq \Sigma \\
&\leq 22 \sum_{j=0}^{m-1} d_{\mathrm{B}}(\gamma(t'_j), \gamma(t'_{j+1})) \text{ by Theorem 3.1 with } K = 1/22 \\
&\leq 22\, L_{d_{\mathrm{B}}}(\gamma).
\end{aligned}$$

[2]Provided the induced metrics are defined using the same class of admissible paths, hence our convention.

Since this is true for any partition Σ of any finite size n, it follows that

$$\hat{d}_{\mathrm{FD}}(\mathrm{R}_f, \mathrm{R}_g) \leq L_{d_{\mathrm{FD}}}(\gamma) \leq 22\, L_{d_{\mathrm{B}}}(\gamma).$$

Again, this inequality is true for any admissible path γ, so $\hat{d}_{\mathrm{FD}}(\mathrm{R}_f, \mathrm{R}_g) \leq 22\, \hat{d}_{\mathrm{B}}(\mathrm{R}_f, \mathrm{R}_g)$. □

Theorem 4.2 implies in particular that $\hat{d}_{\mathrm{B}}$ is a true metric on Reeb graphs, as opposed to d_{B} which is only a pseudo-metric. Moreover, the simplification operator defined in Section 3.1.1 makes it possible to continuously deform any Reeb graph into a trivial segment-shaped graph then into the empty graph. This shows that Reeb is path-connected in d_{FD}. Since the length of such continuous deformations is finite if the Reeb graph is finite, $\hat{d}_{\mathrm{FD}}$ and $\hat{d}_{\mathrm{B}}$ are finite metrics. Finally, the global equivalence of $\hat{d}_{\mathrm{FD}}$ and $\hat{d}_{\mathrm{B}}$ yields the following:

Corollary 4.3. *The metrics $\hat{d}_{\mathrm{FD}}$ and $\hat{d}_{\mathrm{B}}$ induce the same topology on* Reeb*, which is a refinement of the ones induced by d_{FD} or d_{B}.*

Remark 4.4. *Note that the first inequality in (8) and, consequently, Corollary 4.3, are wrong if one defines the admissible paths for $\hat{d}_{\mathrm{B}}$ to be the whole class of maps $[0,1] \to$ Reeb that are continuous in d_{B}—hence our convention. For instance, let us consider the two Reeb graphs R_f and R_g of Figure 1 such that $\mathrm{Dg}(f) = \mathrm{Dg}(g)$, and let us define $\gamma : [0,1] \to$ Reeb by $\gamma(t) = \mathrm{R}_f$ if $t \in [0, 1/2)$ and $\gamma(t) = \mathrm{R}_g$ if $t \in [1/2, 1]$. Then γ is continuous in d_{B} while it is not in d_{FD} at $1/2$ since $d_{\mathrm{FD}}(\mathrm{R}_f, \mathrm{R}_g) > 0$. In this case, $\hat{d}_{\mathrm{B}}(\mathrm{R}_f, \mathrm{R}_g) \leq L_{d_{\mathrm{B}}}(\gamma) = 0 < \hat{d}_{\mathrm{FD}}(\mathrm{R}_f, \mathrm{R}_g)$.*

5 Discussion

In this article, we proved that the bottleneck distance, even though it is only a pseudo-metric on Reeb graphs, can actually discriminate a Reeb graph from the other Reeb graphs in a small enough neighborhood, as efficiently as the other metrics do. This theoretical result legitimates the use of the bottleneck distance to discriminate between Reeb graphs in applications. It also motivates the study of intrinsic metrics, which can potentially shed new light on the structure of the space of Reeb graphs and open the door to new applications where interpolation plays a key part. This work has raised numerous questions, some of which we plan to investigate in the upcoming months:

- **Can the lower bound be improved?** We believe that $\varepsilon/22$ is not optimal. Specifically, a more careful analysis of the simplification operator should allow us to derive a tighter upper bound than the one in Lemma 3.2, and improve the current lower bound on d_{B}.
- **Do shortest paths exist in Reeb?** The existence of shortest paths achieving $\hat{d}_{\mathrm{B}}$ is an important question since a positive answer would enable us to define and study the *intrinsic curvature* of Reeb. Moreover, characterizing and computing these shortest paths would be useful for interpolating between Reeb graphs. The existence of shortest paths is guaranteed if the space is complete and locally compact. Note that Reeb is not complete, as shown by the counter-example of Figure 7. Hence, we plan to restrict the focus to the subspace of Reeb graphs having at most N features with height at most H, for fixed but arbitrary $N, H > 0$. We believe this subspace is complete and locally compact, like its counterpart in the space of persistence diagrams [8].

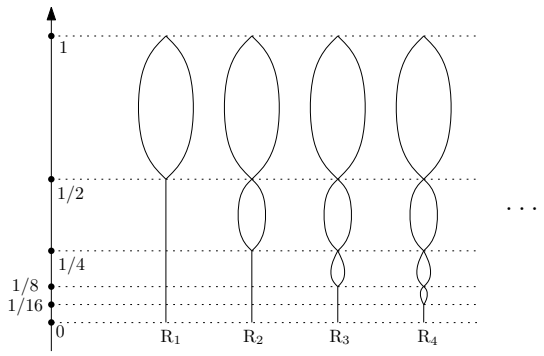

Figure 7: A sequence of Reeb graphs that is Cauchy but that does not converge in Reeb because the number of critical values goes to $+\infty$. Indeed, each R_n has $n+2$ critical values.

- **Is Reeb an Alexandrov space?** Provided shortest paths exist in Reeb (or in some subspace thereof), we plan to determine whether the intrinsic curvature is bounded, either from above or from below. This is interesting because barycenters in metric spaces with bounded curvature enjoy many useful properties [24], and they can be approximated effectively [23].
- **Can the local equivalence be extended to general metric spaces?** We have reasons to believe that our local equivalence result can be used to prove similar results for more general classes of metric spaces than Reeb graphs. If true, this would shed new light on inverse problems in persistence theory.

References

[1] Pankaj Agarwal, Kyle Fox, Abhinandan Nath, Anastasios Sidiropoulos, and Yusu Wang. Computing the Gromov-Hausdorff Distance for Metric Trees. In *Proceedings of the 26th International Symposium on Algorithms and Computation*, 2015.

[2] Muthu Alagappan. From 5 to 13: Redefining the Positions in Basketball. MIT Sloan Sports Analytics Conference, 2012.

[3] Vincent Barra and Silvia Biasotti. 3D Shape Retrieval and Classification using Multiple Kernel Learning on Extended Reeb graphs. *The Visual Computer*, 30(11):1247–1259, 2014.

[4] Ulrich Bauer, Xiaoyin Ge, and Yusu Wang. Measuring Distance Between Reeb Graphs. In *Proceedings of the 30th Symposium on Computational Geometry*, pages 464–473, 2014.

[5] Ulrich Bauer, Elizabeth Munch, and Yusu Wang. Strong Equivalence of the Interleaving and Functional Distortion Metrics for Reeb Graphs. In *Proceedings of the 31st Symposium on Computational Geometry*, 2015.

[6] Silvia Biasotti, Daniela Giorgi, Michela Spagnuolo, and Bianca Falcidieno. Reeb Graphs for Shape Analysis and Applications. *Theoretical Computer Science*, 392:5–22, 2008.

[7] Håvard Bjerkevik. Stability of higher-dimensional interval decomposable persistence modules. *CoRR*, abs/1609.02086, 2016.

[8] Andrew Blumberg, Itamar Gall, Michael Mandell, and Matthew Pancia. Robust Statistics, Hypothesis Testing, and Confidence Intervals for Persistent Homology on Metric Measure Spaces. *Foundations of Computational Mathematics*, 14:745–789, 2014.

[9] Dmitri Burago, Yuri Burago, and Sergei Ivanov. *A Course in Metric Geometry*. American Mathematical Society, 2001.

[10] Gunnar Carlsson, Vin de Silva, and Dmitriy Morozov. Zigzag Persistent Homology and Real-valued Functions. In *Proceedings of the 25th Symposium on Computational Geometry*, pages 247–256, 2009.

[11] Mathieu Carrière and Steve Oudot. Structure and Stability of the 1-Dimensional Mapper. In *Proceedings of the 32nd Symposium on Computational Geometry*, volume 51, pages 25:1–25:16, 2016.

[12] Frédéric Chazal, Vin de Silva, Marc Glisse, and Steve Oudot. *The Structure and Stability of Persistence Modules*. Springer, 2016.

[13] David Cohen-Steiner, Herbert Edelsbrunner, and John Harer. Extending persistence using Poincaré and Lefschetz duality. *Foundation of Computational Mathematics*, 9(1):79–103, 2009.

[14] Vin de Silva, Elizabeth Munch, and Amit Patel. Categorified Reeb Graphs. *Discrete and Computational Geometry*, 55:854–906, 2016.

[15] Barbara di Fabio and Claudia Landi. The Edit Distance for Reeb Graphs of Surfaces. *Discrete and Computational Geometry*, 55(2):423–461, 2016.

[16] Marcio Gameiro, Yasuaki Hiraoka, and Ippei Obayashi. Continuation of Point Clouds via Persistence Diagrams. *Physica D: Nonlinear Phenomena*, 334:118–132, 2016.

[17] Xiaoyin Ge, Issam Safa, Mikhail Belkin, and Yusu Wang. Data Skeletonization via Reeb Graphs. In *Advances in Neural Information Processing Systems 24*, pages 837–845, 2011.

[18] Alexandr Ivanov, Nadezhda Nikolaeva, and Alexey Tuzhilin. The Gromov-Hausdorff Metric on the Space of Compact Metric Spaces is Strictly Intrinsic. *Mathematical Notes*, 100(6):947–950, 2016.

[19] P. Lum, G. Singh, A. Lehman, T. Ishkanov, M. Vejdemo-Johansson, M. Alagappan, J. Carlsson, and G. Carlsson. Extracting insights from the shape of complex data using topology. *Scientific Reports*, 3, 2013.

[20] Waleed Mohamed and A. Ben Hamza. Reeb graph path dissimilarity for 3d object matching and retrieval. *The Visual Computer*, 28(3):305–318, 2012.

[21] Tomoyuki Mukasa, Shohei Nobuhara, Atsuto Maki, and Takashi Matsuyama. Finding articulated body in time-series volume data. In *Proceedings of the 4th International Conference on Articulated Motion and Deformable Objects*, pages 395–404, 2006.

[22] Monica Nicolau, Arnold Levine, and Gunnar Carlsson. Topology based data analysis identifies a subgroup of breast cancers with a unique mutational profile and excellent survival. *Proceedings of the National Academy of Science*, 108(17):7265–7270, 2011.

[23] Shin-Ichi Ohta. Gradient flows on Wasserstein spaces over compact Alexandrov spaces. *American Journal Mathematics*, 131(2):475–516, 2009.

[24] Shin-Ichi Ohta. Barycenters in Alexandrov spaces of curvature bounded below. *Advances in Geometry*, 12:571–587, 2012.

[25] Georges Reeb. Sur les points singuliers d'une forme de Pfaff complètement intégrable ou d'une fonction numérique. *Compte Rendu de l'Académie des Science de Paris*, 222:847–849, 1946.

[26] Gurjeet Singh, Facundo Mémoli, and Gunnar Carlsson. Topological Methods for the Analysis of High Dimensional Data Sets and 3D Object Recognition. In *Symposium on Point Based Graphics*, 2007.

[27] Julien Tierny, Jean-Philippe Vandeborre, and Mohamed Daoudi. Invariant High Level Reeb Graphs of 3D Polygonal Meshes. *International Symposium on 3D Data Processing Visualization and Transmission*, pages 105–112, 2006.